Technical white paper

Data Services Cloud Console Security Guide

Contents

Executive summary

Cloud-based management from Hewlett Packard Enterprise offers many advantages for both data infrastructure and the data itself. Data Services Cloud Console (DSCC) is the HPE cloud-based application for current and future data and storage management. Security is at the center of the design of DSCC. This paper describes how data security, from on-premises data centers to cloud management, is built into the application. It describes how HPE developed DSCC and its associated products to meet customers' security needs. It also details how DSCC provides a secure platform on which customer arrays and the data they contain can be securely accessed and managed.

This paper is designed for corporate security personnel, risk planners, and storage array administrative personnel who will be working with DSCC. Readers should be familiar with computer concepts associated with management, networking, security, and data storage arrays, as well as with the needs and requirements of their organization and its infrastructure.

After reading this document, customers should understand how the security capabilities of DSCC are enabled and implemented.

Overview

DSCC is an application that is fully integrated into the HPE GreenLake edge-to-cloud platform, leveraging security features and resources that are common to other applications.

HPE GreenLake edge-to-cloud platform

The HPE GreenLake edge-to-cloud platform offers a common set of cloud services that enable a consistent, cloud-qualified customer experience. The HPE GreenLake platform is designed to combine the cloud's agility with the governance, compliance, and visibility of the hybrid IT model.

Key features of HPE GreenLake make it easy for new cloud users to get started while offering powerful capabilities for advanced users:

- Global data management for streamlined configuration and deployment of devices. HPE GreenLake supports device management, which enables customers to provision and manage multiple devices that have similar configuration requirements with less administrative overhead.
- Secure cloud-based platform.
- Rich API that enables customers to implement data management functionality.

Data Services Cloud Console

DSCC is a secure cloud application, deployed on HPE GreenLake, that provides a control plane for simplifying data infrastructure management and delivering data services across edge-to-cloud environments, as shown in Figure 1. This enables a unified management experience for the customer's cloud-enabled arrays—their fleet.

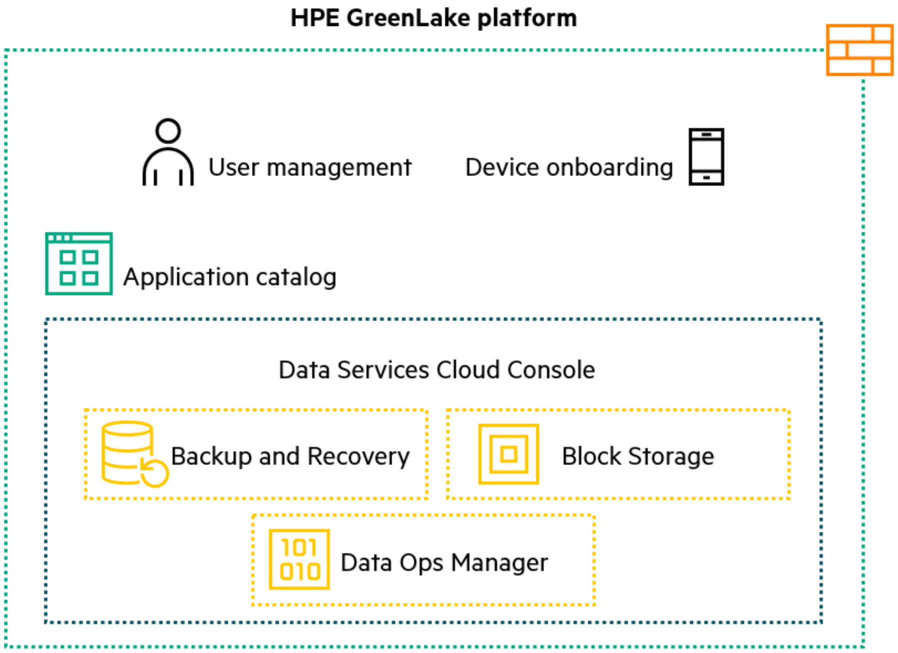

Figure 1. DSCC services overview

DSCC contains several software-as-a-service (SaaS) offerings:

- **Backup and Recovery:** Enables automated data protection for VMware® resources by using prebuilt and customizable protection policies to create app-consistent, immutable recovery points.
- **Block Storage:** Enables simple provisioning and protection of block storage based on the application's required storage tier, with intelligent recommendations that can help ensure uniform provisioning and utilization across all arrays.
- **Data Ops Manager:** Enables customers to monitor all their systems from a single dashboard, create fleet-wide host definitions, and view and apply software updates for their entire install base with federated updates.

Note

For additional information on DSCC, refer to the DSCC technical white paper.

Keeping customer's data secure

The main objective of the security design and features of DSCC is to ensure that a customer's data can be always kept secure. To achieve this goal, it is important to understand the types of data, how and where they are stored, the responsibilities, and how the data is accessed.

Data collection and processing

Data collection and processing, together with data residency and sovereignty, is a growing concern for users of SaaS solutions. Hewlett Packard Enterprise addresses these issues in the following ways:

- Hewlett Packard Enterprise has established the Global Privacy Program to ensure that its processing of personal data complies with applicable laws and regulations, including ensuring that Hewlett Packard Enterprise follows the approach of privacy and security by design for its products and services.
- HPE GreenLake stores user and organizational account information in the US, adhering to applicable privacy and security laws and regulations. Hewlett Packard Enterprise acts as a data controller for the processing of user and organizational account information.
- DSCC application data is stored at the regional level. DSCC data collection is limited strictly to configuration and performance-related data, as shown in Table 1. HPE acts as a data processor for any personal data it may process on behalf of a customer as part of DSCC application.
- Data written to the on-premises devices (for example, data stored within array volumes or virtual machines) is not exposed to or accessible from DSCC.

The information retained by HPE GreenLake and DSCC consists of the information listed in Table 1.

Table 1. Data stored in HPE GreenLake and DSCC

Location	Information	Details stored
HPE GreenLake	User details	• Name, email address, and username
HPE GreenLake	Organizational details	• Organization name, address, and contact details
DSCC	Device inventory	• Product family and model • System details, including health status, performance, and performance policies • Enclosure information: Disks, cards, fans, expanders, and power • Node information: Disks, cards, CPUs, MCUs, memory, batteries • Storage information: Capacity usage, virtual volume and vLUN information, including host and host sets, capacity utilization, volume collections, and snapshots • Backup and protection: Data management jobs and templates, protection schedules and templates, and policies • Network and ports: Configurations, interfaces, port settings, proxy settings, and service ports • Fibre Channel configuration, interface, port, and initiator details • Software: Version, history, update status, and update recommendations • Virtualization: Hypervisor details (folders, networks, resource pools, hosts, clusters), VM information, datastore information, snapshot information • VMware vCenter Server® settings, and collections • Support settings • Monitoring and reporting: Alerts, tasks, alarms, SNMP details, mail settings, and audit logs

Shared security control

Both HPE and the customer have a part to play in controlling security—either in the cloud, where Hewlett Packard Enterprise takes measures to prevent unauthorized access to user profile data, or at the customer site, where the customer prevents unauthorized access to IT hardware.

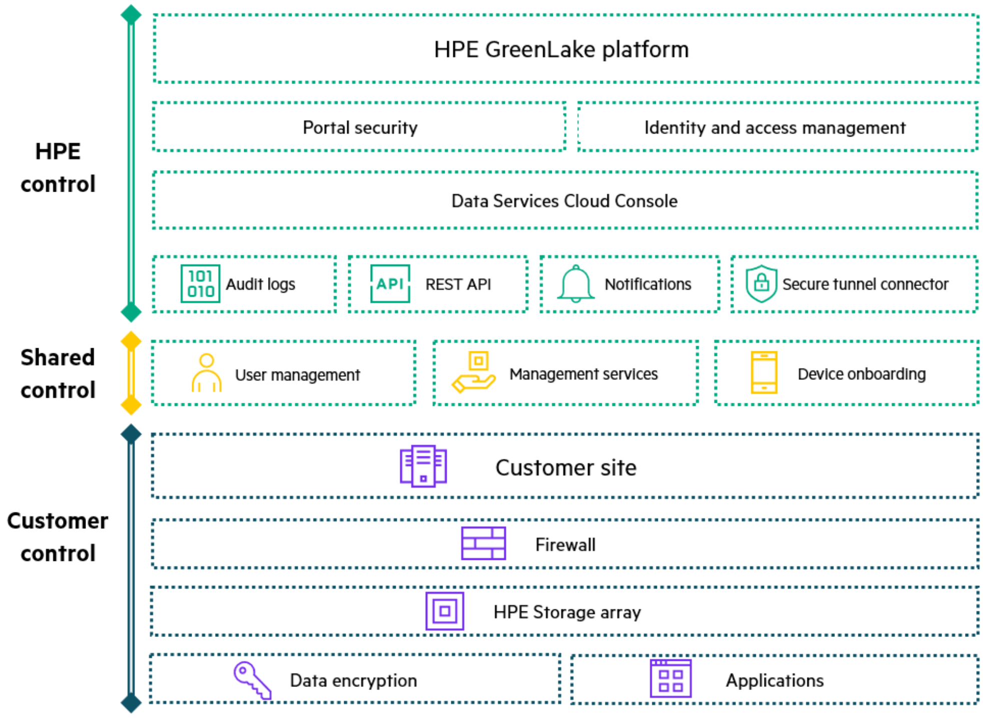

Figure 2. DSCC shared security control

Hewlett Packard Enterprise follows best practices for SaaS hosting, which include the use of cloud security tools (firewalls, distributed load balancers, monitoring, and auditing). Both the customer and Hewlett Packard Enterprise share control over the securing of user credentials for accessing HPE GreenLake and DSCC. Customers are responsible for securing their premises, managing access to the company's networks, restricting physical access to the HPE array containing encrypted volumes hosting customer applications, and implementing recommended best practices.

In the case of a support engagement with Hewlett Packard Enterprise, HPE service personnel are only able to access the on-premises devices after the customer provides explicit approval. This restriction applies to remote sessions such as a screen-sharing, as well as physical access to the device on the customer's premises and to access through DSCC. Access through DSCC management features is possible only after the account administrator has explicitly granted authorization to a user, and the account administrator can revoke access at any time.

Access to data

Data that is written to the on-premises arrays and on-premises components, such as virtual machines, is not exposed to DSCC. Data stored on the device stays on the device. Although the array requires a secure connection to DSCC, it is not exposed to the internet.

Only authorized users within the customer's DSCC account can send API commands to the on-premises arrays. HPE engineers and service personnel can only perform such actions from DSCC with expressed customer permission. Every action executed in DSCC is recorded in the audit logs.

HPE GreenLake and DSCC are maintained by the HPE operations team, which has no access to customer account information. The update process involves deploying the latest versions of software and managing the monitoring and auditing functions. No one else has access to the production instance to effect any changes—not even engineers working on the software that is deployed in production. As part of the secure software development lifecycle, any software change, whether it is new software or an update, must be peer-reviewed and tested before it can be deployed.

Note

Hewlett Packard Enterprise conducts employment screenings (background checks) on all new external hires, where legally permissible.

Device setup and connection to the cloud

The customer's first task is to set up a new user account in HPE GreenLake. The next task is to create an organization account. This organization account in HPE GreenLake, also known as an **organization unit** or **company account**, is the single account that all users and devices from a single organization are associated with. After the organization account is created, the account administrator can onboard the storage devices to it.

On-premises devices

To manage their fleet, customers deploy two types of devices:

- **Hardware**: HPE arrays located at the customer site, such as HPE Alletra and HPE Primera.
- **Software**: Virtual machines deployed on an on-premises hypervisor to provide specific local functions for DSCC services, such as Data Orchestrator and Protection Store Gateway (PSG), used by the HPE Backup and Recovery service.

User and site public interfaces

To connect to HPE GreenLake and DSCC from a web browser, a client application using the API, or HPE arrays, the user must have access to the public interfaces of DSCC to use the services provided.

Note

The customer's network firewall must allow port 443/tcp (HTTPS) outbound connections to these public interfaces. No inbound connection is required.

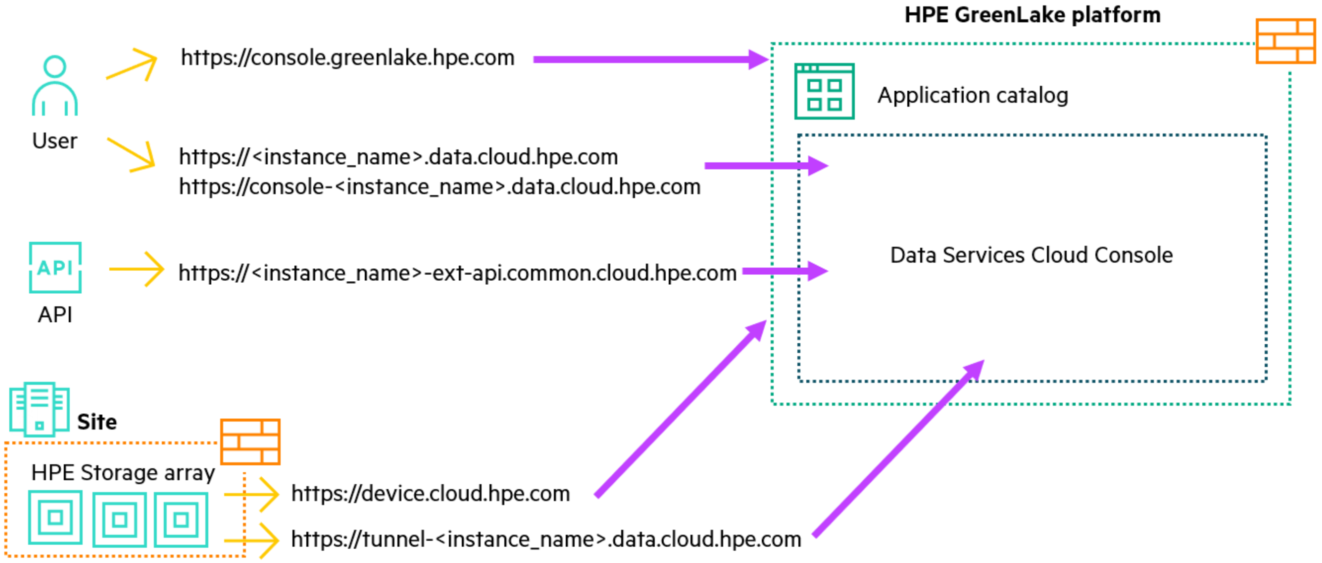
HPE GreenLake platform
User
https://console.greenlake.hpe.com
https://<instance_name>.data.cloud.hpe.com
https://console-<instance_name>.data.cloud.hpe.com
API
API
https://<instance_name>-ext-api.common.cloud.hpe.com
Application catalog
Data Services Cloud Console
Site
HPE Storage array
https://device.cloud.hpe.com
https://tunnel-<instance_name>.data.cloud.hpe.com

Figure 3. DSCC public interfaces

Table 2. DSCC public interfaces details

Public interface FQDN	Port	Initiator	Description
https://console.greenlake.hpe.com	443/tcp	User	Allows communication to HPE GreenLake
https://console-<instance name>.data.cloud.hpe.com	443/tcp	User	Allows communication to DSCC instance (America: **us1**, Asia: **jp1**, Europe: **eu1**) For example, **https://console-us1.data.cloud.hpe.com**
https://device.cloud.hpe.com	443/tcp	HPE array	Required for hardware device activation
https://tunnel-<instance name>.data.cloud.hpe.com	443/tcp	HPE array, Data Orchestrator	Allows communication to DSCC instance (us1, jp1, eu1) For example, **https://tunnel-jp1.data.cloud.hpe.com**
https://<instance name>.data.cloud.hpe.com	443/tcp	User/API	Allows communication to DSCC for user and API access (us1, jp1, eu1) For example, **https://eu1.data.cloud.hpe.com**

Public interface FQDN	Port	Initiator	Description
https://<instance name>-ext-api.common.cloud.hpe.com	443/tcp	API	Allows communication to DSCC for API access (us1, jp1, eu1) For example, **https://eu1-ext-api.common.cloud.hpe.com**
https://cosm-*.s3.*.amazonaws.com	443/tcp	PSG	Allows communication to AWS S3 buckets
https://midway.ext.hpe.com	443/tcp	Data Orchestrator	Required for software device activation and connection to HPE Remote Device Access (RDA) and HPE InfoSight

Note
HPE arrays support the use of network proxies to access DSCC public interfaces.

Network load balancers are used to prioritize all traffic to DSCC instances. In the event of a communications disconnect between DSCC and the HPE array, the network load balancer automatically begins the process of re-establishing the connection to DSCC by generating a new session key. After the secure tunnel is re-established, any queued events since the disconnect are transmitted.

Hardware device onboarding and activation

HPE arrays must be onboarded with HPE GreenLake before DSCC can manage them. After the onboarding process is complete, the HPE array creates a secure mutual Transport Layer Security (mTLS) tunnel to HPE GreenLake for activation.

All HPE arrays contain a per-device unique client certificate signed by a trusted certificate authority (CA) and installed during the manufacturing process. This certificate is used when setting up a connection to uniquely identify the device to HPE GreenLake and DSCC, which validate the client certificate by using its authoritative entity (trust anchor).

Important
The onboarding process requires the serial number of the HPE array and the DSCC subscription key to be registered within HPE GreenLake.

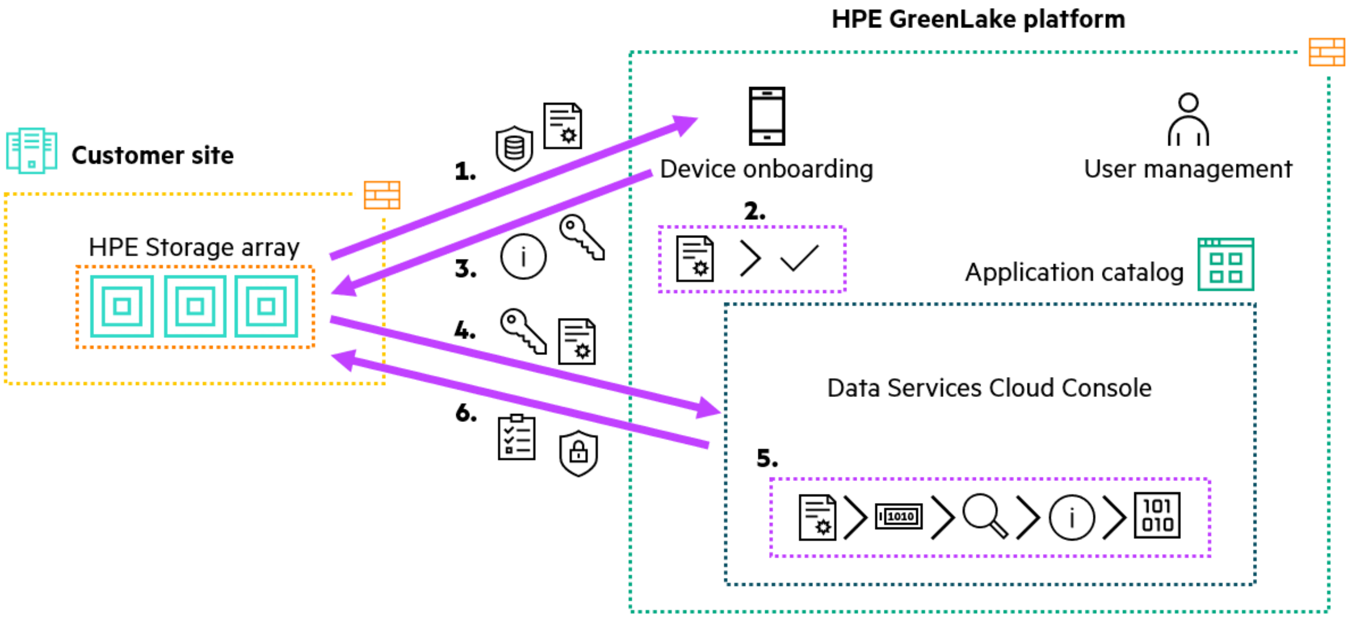

Figure 4. HPE array onboarding and activation process

The hardware device onboarding and activation process includes the following steps:

1. The on-premises array connects to HPE GreenLake through a secure tunnel using its client certificate.
2. HPE GreenLake validates the client certificate by using its authoritative entity. HPE GreenLake uses the information in the certificate to identify the serial number of the array. The customer selects the DSCC instance (region) to which the array will connect.
3. HPE GreenLake returns the address of the DSCC instance to be used by the array and provides a subscription key that is readable by the array software.
4. The array connects directly to the selected DSCC instance through a secure tunnel, using its client certificate. The array keeps trying to connect with DSCC until a secure connection can be made. When a secure connection is successfully established, the DSCC instance uses its authoritative entity to validate the certificate of the array.
5. DSCC extracts the serial number of the array from the supplied certificate, looks up the tenant ID of the customer's organization, and—after a match is found—allows activation of the device ID against the serial number.
6. DSCC returns an ID to the array, and the array adds it to the trusted list.

Software device deployment and activation

HPE virtual machines that connect to DSCC are downloaded as generic images and must be deployed in a local hypervisor. During deployment, they are assigned a unique identity and a client certificate that validates the device's identity and associates it with a DSCC instance and a customer account.

As part of the HPE Backup and Recovery service, two virtual machines are deployed on premises:

- Data Orchestrator maintains secure connectivity from the on-premises environment to the cloud services. It is responsible for governing the schedules locally, maintaining and orchestrating the workflows and tasks across several components such as the vCenter Servers, VMware ESXi™ servers, arrays, and PSGs.
- The PSG is the backup target built using HPE Catalyst protocol to deliver unmatched storage efficiency for the HPE Backup and Recovery service.

Only Data Orchestrator establishes a secure tunnel with DSCC—the PSG virtual machine never connects to DSCC. If customers choose to copy their data to an HPE Backup and Recovery service Cloud Protection Store, PSG will always send the data over an encrypted link.

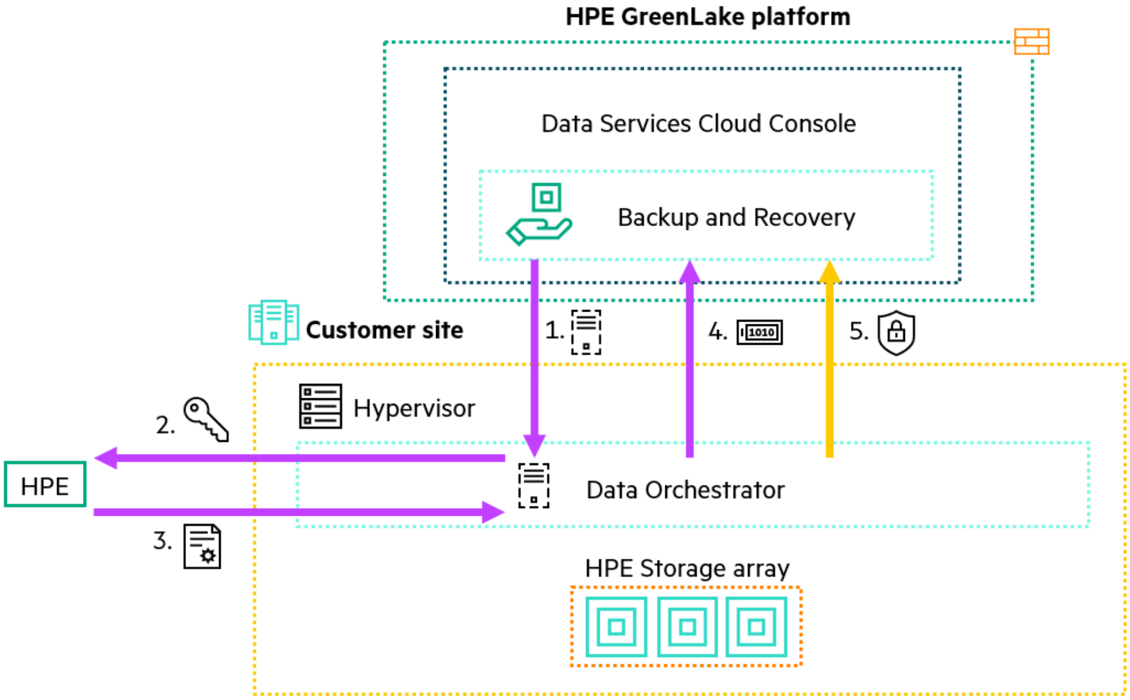

Figure 5. Software device deployment and activation process

The process of deploying and activating the Data Orchestrator virtual machine that is part of the HPE Backup and Recovery service follows these steps:

1. A user with the appropriate role downloads the Data Orchestrator image and deploys it to a on-premises hypervisor. As part of the deployment, the user configures the network and proxy settings to allow the virtual machine to connect to web services.
2. Data Orchestrator generates a local private key and sends a certificate-signing request to Hewlett Packard Enterprise.
3. A client certificate is generated from the received certificate-signing request and sent to Data Orchestrator. The device keeps trying to connect to DSCC until Step 4 is completed.
4. An activation code is generated on Data Orchestrator, and the user must manually assign the code to an HPE Backup and Recovery service application within its DSCC instance. After the assigned activation code is validated, it becomes a unique identifier for the virtual machine.
5. After the virtual machine's unique identifier is assigned to a DSCC instance, the secure tunnel is established.

Important
The request for a secure tunnel is always initiated by the on-premises device and never by DSCC. It is not possible for DSCC to initiate a connection to an on-premises device.

Management tools

HPE arrays enable customers to take advantage of different management tools to better suit their environment and requirements through the management GUI, CLI, RESTful APIs, or plug-ins. DSCC enhances the customer experience, providing unified fleet management. To enforce user security on DSCC, users created on the HPE array do not automatically gain access or equivalent permissions within DSCC.

Important
Customer data stored on HPE array volumes or LUNs cannot be accessed by and is never sent to DSCC. It is not possible to open a CLI or shell session from DSCC to the array. Only API commands are sent from DSCC to the array.

Secure sites

Customer secure sites typically contain systems that cannot communicate with other devices outside the customer's internal network. On-premises arrays must have a connection to HPE GreenLake for activation and to DSCC for software and firmware updates.

Identity and access management

This section describes the methods of authenticating and authorizing access to DSCC.

Authentication

Access to DSCC requires a unique HPE GreenLake user credential. This credential is linked to the user's name and email address through the registration process in HPE GreenLake. After the user is registered, the submitted data is verified to ensure that it meets global trade regulations. After verification, the user account must be joined to the organization account.

After the software subscription is activated, users must create both their own user account in HPE GreenLake and an organization account. This organization account or company account in HPE GreenLake is the single account that all users and devices from a single organization are associated with. The organization account has at least one user with the account administrator role. The creator of an organization account is automatically assigned the account administrator role with administrator privileges for that organization and has the option of inviting additional users and configuring access permissions.

Figure 6 shows the capabilities and options that HPE GreenLake offers customers to enable them to customize their authentication preferences to securely manage their HPE array through the DSCC application services. As the diagram shows, the connections between DSCC and the HPE array are only initiated by the HPE array.

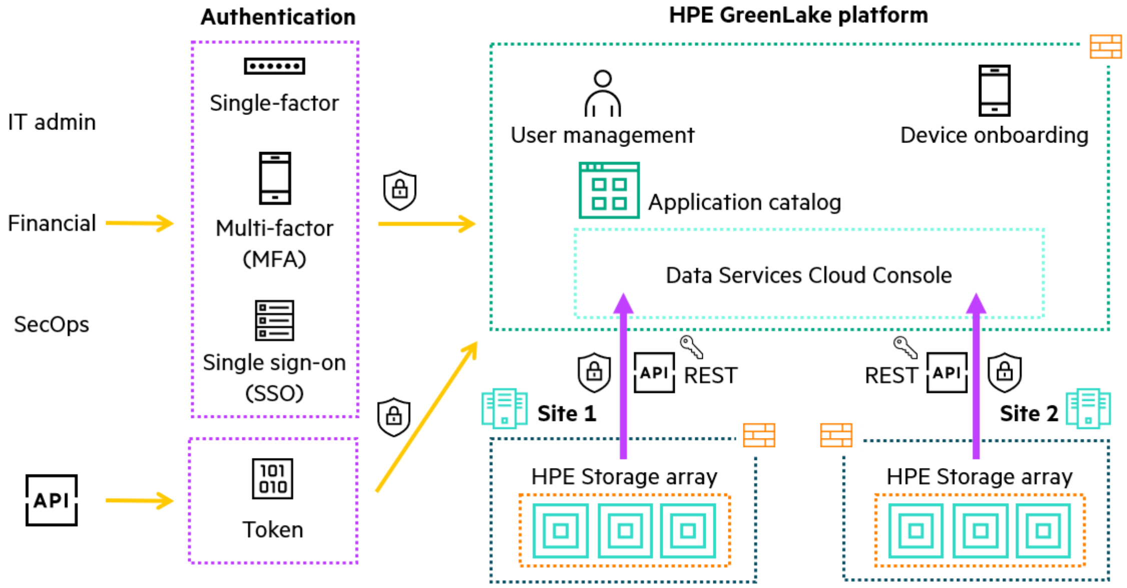

Figure 6. HPE GreenLake authentication overview

Note

All user account and organizational information created in HPE GreenLake is stored in the US. For more information about how personal data is protected, see the HPE Privacy Statement.

A single user account can create or join multiple organizational units. Because HPE GreenLake is a multi-tenant environment, it allows users to access only the information that belongs to their assigned organizational units.

If a user no longer requires access to HPE GreenLake and DSCC, the administrator can disable the user's account. For traceability and compliance, audit logs referencing the user remain unchanged to accurately reflect actions performed by that user.

Important

Hewlett Packard Enterprise recommends that organizations include the removal of HPE GreenLake user accounts in their employee offboarding process.

A user can log in to HPE GreenLake by:

- Creating and managing an account in HPE GreenLake
- Using single sign-on (SSO)

Creating user accounts in HPE GreenLake

User accounts created in HPE GreenLake provide access to all HPE GreenLake apps and services, depending on the level of authorization the user has been assigned. The account credentials are unique to HPE GreenLake.

To create user accounts, an administrator adds new users to the organization by sending an email invitation from HPE GreenLake. The email includes a link to create a new user account. This link contains a randomly generated value that is unique to the email address to which it was sent; it cannot be used to create an account in HPE GreenLake with a different email address. When adding the user, the administrator also specifies the role that the user will be assigned after creating an account and logging in for the first time.

User accounts created in HPE GreenLake can use multifactor authentication (MFA) for additional security.

Important

Hewlett Packard Enterprise strongly recommends that customers take advantage of MFA to protect their accounts.

HPE GreenLake account passwords

Passwords must meet the following specifications:

- Must be between eight and 255 characters, with a minimum of five unique characters
- Cannot have more than two repeated characters
- Must contain at least one special character, one numeric character, one uppercase letter, and a lowercase letter
- Cannot contain user account data and are checked against a list of commonly known passwords

Also, all validated passwords must adhere to a strict policy:

- New passwords cannot match the past six passwords used within the past 365 days.
- Passwords expire automatically after 182 days.
- Accounts with multiple login failures are automatically locked out for a period.

Multifactor authentication

MFA implements an additional level of authentication for a user to gain access to HPE GreenLake.

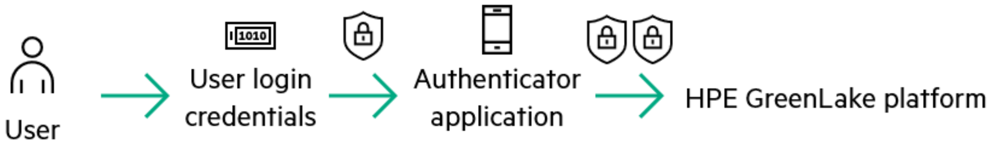

Figure 7. Software-based authenticator login process

Within HPE GreenLake, MFA can be enabled on two levels:

- **Organization account level**. This option enforces the use of MFA for all user accounts within the organization created in HPE GreenLake. It is configured in the organization account settings by an administrator user with the appropriate role, and it cannot be disabled by an individual user. When this option is enabled, users are forced to configure MFA either at their next login or when they create a new user account.
- **Individual user account level**. If MFA is not enabled at the organization level by an account administrator, individual users can enable MFA for their own accounts in HPE GreenLake with no impact to any other user account configured within the organization in HPE GreenLake.

HPE GreenLake supports standard software-based authenticators. As part of the setup process of MFA, a QR code is displayed in the HPE GreenLake UI that the user can scan by using an MFA app.

If a user is a member of multiple organizations in HPE GreenLake, and any one of those organizations enforces the use of MFA, then the user must set up MFA for their account the next time they log in to HPE GreenLake.

Single sign-on authentication

SSO authentication enables organizations to simplify the user experience of logging in to external portals by allowing users to enter their company login credentials to authenticate their identity. HPE GreenLake uses Security Assertion Markup Language (SAML) 2.0 as the standard protocol for SSO.

Users can log in to HPE GreenLake with their company user credentials. HPE GreenLake redirects the user to the identify provider's login page, which verifies the user's credentials and, if successful, redirects them back to HPE GreenLake, where they will be logged in to their account.

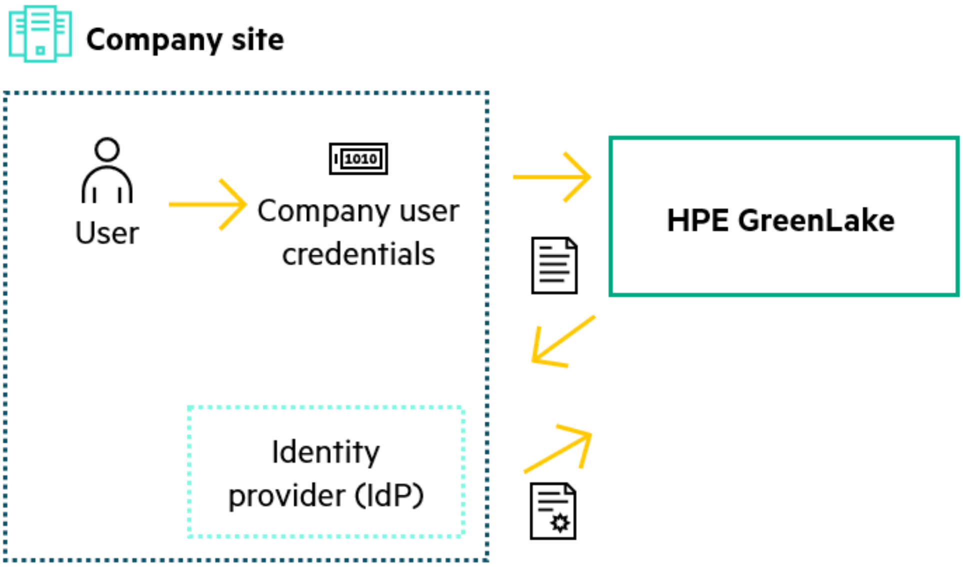

Figure 8. HPE GreenLake SSO integration process

Note

With SSO, an HPE GreenLake account is not required if users are invited by the organizational unit administrator, allowing them instant access into HPE GreenLake with their company user credentials. If a user account is disabled in their organization, access to HPE GreenLake is automatically revoked.

Users who log in with SSO have an HPE GreenLake account created automatically when they log in for the first time, and they are audited in the same way as users accounts created in HPE GreenLake.

Important

MFA that is configured in HPE GreenLake is applicable only to user accounts that have been configured in HPE GreenLake. If the customer is using SSO, then any MFA configuration is managed by the customer's identity provider (IdP) and invited users do not require a HPE GreenLake user account to be created.

Configuring SAML SSO for HPE GreenLake

The HPE GreenLake SAML 2.0 SSO consists of the key elements listed in Table 3.

Table 3. Single sign-on SAML 2.0 attributes

Element	Description
Service provider (SP)	The provider of a business function or service: HPE GreenLake. The SP requests and obtains an identity assertion from the IdP. Based on this assertion, the SP allows a user to access the service.
Identity provider (IdP)	The identity management system maintains the user's identity information and authenticates the user.
SAML request	The authentication request is generated when a user tries to access HPE GreenLake.
SAML assertion	The authentication and authorization information issued by the IdP allows access to the HPE GreenLake service.
Relying party	The relying party is the business service that relies on SAML assertion for authenticating a user (for example, HPE GreenLake).
Asserting party	The asserting party is the Identity management system or the IdP that creates SAML assertions for an SP.
Metadata	Data in the XML format is exchanged between the trusted partners (IdP and HPE GreenLake) to establish interoperability.

Element	Description
SAML attributes	SAML attributes are the attributes associated with the user (for example, username, customer ID, role, and group in which the devices belonging to a user account are provisioned). The SAML attributes must be configured on the IdP according to specifications associated with a user account in HPE GreenLake. These attributes are included in the SAML assertion when HPE GreenLake sends a SAML request to the IdP.
Entity ID	The entity ID is a unique string to identify the service provider that issues a SAML SSO request.
User	In this context, a user has SSO credentials.

HPE GreenLake supports SP-initiated SSO. In an SP-initiated workflow, the SSO request originates from the SP domain. The SSO configuration includes the following steps:

1. Configure the SSO profile in HPE GreenLake.
2. Configure IdP metadata such as entity ID, login URL, logout URL, and X.509 signing certificate.
3. Specify the SAML 2.0 attributes to be associated with a user account.
4. Create a recovery user.

Details that must be provided when setting up the customer's IdP SSO are listed in Table 4.

Table 4. SSO SAML 2.0 settings

SAML setting name	Value
EntityID	**https://sso.common.cloud.hpe.com**
AssertionConsumerService (login URL)	**https://sso.common.cloud.hpe.com/sp/ACS.saml2**
SingleLogoutService	**https://sso.common.cloud.hpe.com/sp/SLO.saml2**

In addition to the SAML 2.0 metadata, the SAML attributes listed in Table 5 must be associated with the user account.

Table 5. SAML 2.0 user attributes

Attribute name	Value
hpeGreenLake	user.hpe_ccs_attribute
NameId	user.email

When `hpe_ccs_attribute` is assigned to a user, the attribute must be in the following format:

```
version_1#<platform_customer_id>:<GreenLake ID>:<GreenLake role>:ALL_SCOPES:<application
ID>:<application role>:ALL_SCOPES
```

The information required in this attribute string is in the HPE GreenLake UI under **Manage account** → **Single Sign On (SSO)** → **View SAML attribute**.

Authorization

The authorization of users can be controlled by assigning user permissions or by assigning roles.

User permissions

HPE GreenLake user permissions are enforced by using role-based access control (RBAC) to ensure that the correct level of access is given to each user. The administrator of an organizational unit has the option to assign predefined roles provided by HPE GreenLake or to create custom roles for users.

Note

The creator of the organizational unit within HPE GreenLake is automatically assigned administrator permissions.

Table 6. HPE GreenLake permissions terminology

Type	Assigned permissions
User	An HPE GreenLake user account
Assignment	A list of roles and scopes
Role	A list of permissions
Permission	Create, edit, and delete functionality, among others
Scope	A list of resources that are affected by the roles assigned (for example, only resources in the EU region can be edited)

Although customers can create their own roles, HPE GreenLake offers a few predefined examples, as listed in Table 7.

Table 7. HPE GreenLake role examples

Role	Assigned permissions
Account administrator	Is administrator for the HPE GreenLake account. All permissions are available to set up and configure HPE GreenLake settings and application-level actions (array functions).
Operator	Can view and modify settings.
Observer	Has view-only access, cannot execute any action.

Each application in DSCC has its own set of predefined roles, as listed in Table 8.

Table 8. Predefined DSCC roles

Role	Assigned permissions
Administrator	All permissions available, including storage system and application-level actions.
Read only	Has view-only permissions, cannot apply any changes.
Data Ops Manager administrator	Has all Data Ops Manager permissions.
Data Ops Manager operator	Has Data Ops Manager permissions to create and edit functions, but not delete.
Backup and Recovery administrator	Has all Backup and Recovery administration permissions, including protection policies.
Backup and Recovery operator	Has Backup and Recovery permissions except for application registration and to create or delete protection policies.

Role assignment

HPE GreenLake role assignments offer a diverse selection of permission options available to the administrator. In the example illustrated in Figure 9, the administrator can associate multiple assignments with a single user, with each assignment consisting of a range of roles, such as systems administrator (a predefined collection of permissions), and scopes, such as EU region (only European resources are accessible).

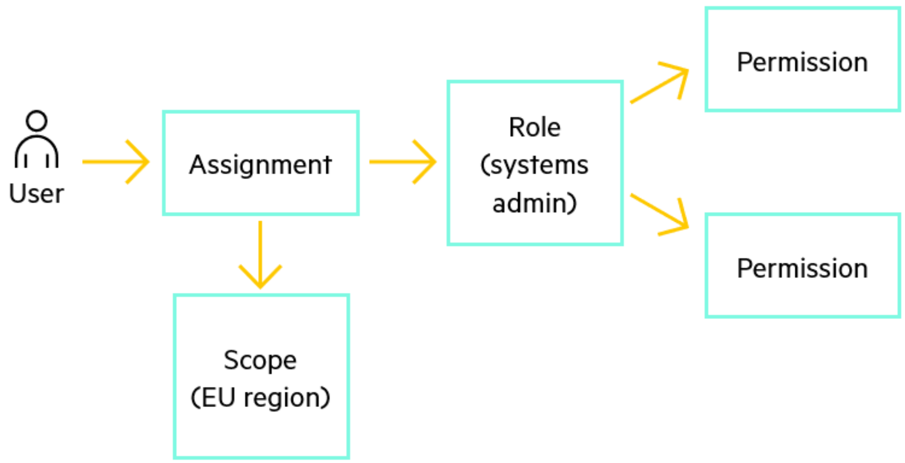

Figure 9. Role assignment process

Note

Only engineers authorized by Hewlett Packard Enterprise who have customer consent have access to customer accounts for support troubleshooting. The customer can revoke such access at any time.

Resource Restriction Policies

The Resource Restriction Policies feature enables the creation of customizable named lists of resources that can be used to limit access to the scope of resources. This feature replaces the need to specify multiple individual resource scopes. Resource scopes are defined by the application and are added to a user role assignment where permissions apply. Creating resource restriction policies allows HPE GreenLake administrators to reuse the same list of resource scopes across multiple assignments, simplifying operations by removing the possibility of missing scopes across multiple assignments for multiple users.

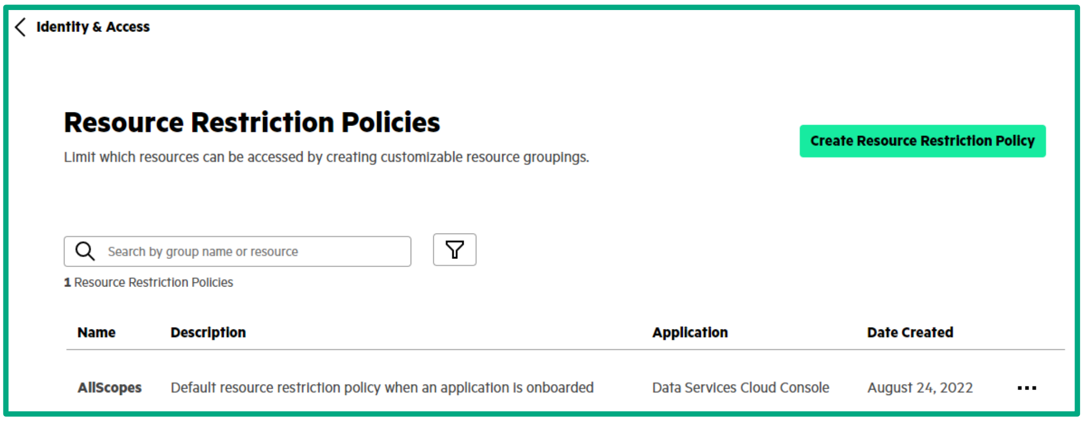
Identity & Access
Create Resource Restriction Policy
Search by group name or resource

Resource Restriction Policies

Limit which resources can be accessed by creating customizable resource groupings.

1 Resource Restriction Policies

Name	Description	Application	Date Created
AllScopes	Default resource restriction policy when an application is onboarded	Data Services Cloud Console	August 24, 2022

Figure 10. HPE GreenLake Resource Restriction Policies screen

For more information on how to set up resource restriction policies, refer to the HPE GreenLake edge-to-cloud platform user guide.

IP Access Rules

The administrator of an organizational unit can enable IP Access Rules in HPE GreenLake to further increase the level of control of who can access HPE GreenLake services. This IP-based access control management resource enables the restriction of access to only pre-authorized IP addresses, networks, or ranges. As an optional access control feature, IP Access Rules work with the standard authentication and authorization features of HPE GreenLake.

After an administrator enables IP Access Rules, only users accessing from an IP address within the specified list of access rules can gain access to their HPE GreenLake account and manage DSCC.

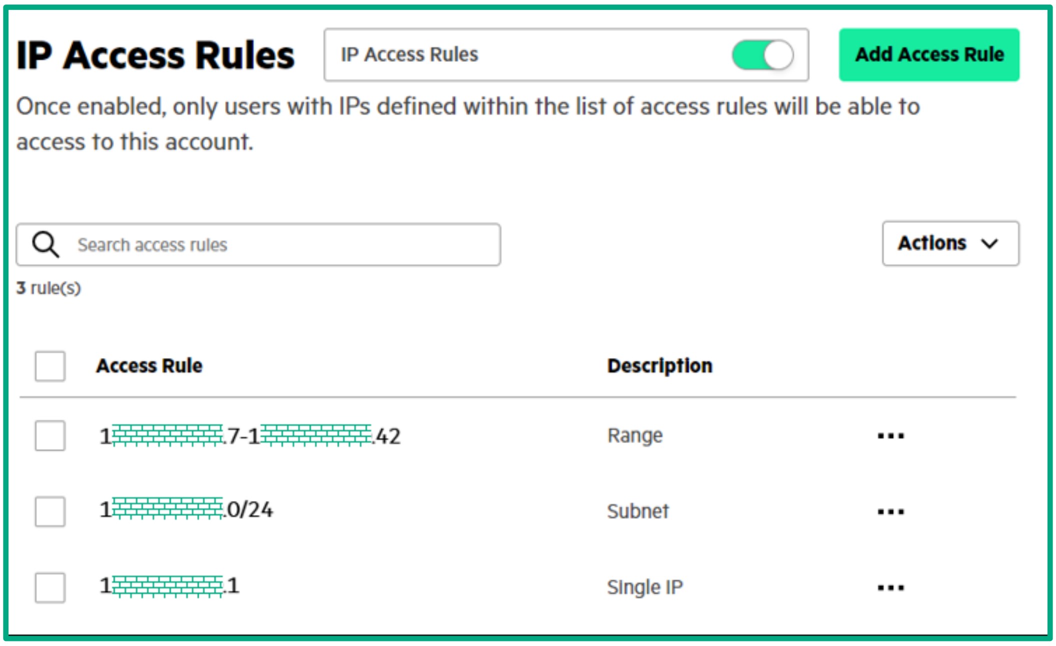

IP Access Rules

Once enabled, only users with IPs defined within the list of access rules will be able to access to this account.

3 rule(s)

□	Access Rule	Description
□	1[redacted].7-1[redacted].42	Range
□	1[redacted].0/24	Subnet
□	1[redacted].1	Single IP

Figure 11. HPE GreenLake IP Access Rules screen

Important

If the external IP address, network, or range used by the administrator when activating IP Access Rules is not included in the initial list, the user will be locked out of HPE GreenLake. To prevent this, a warning is displayed, and the administrator must enter ENABLE before confirming the action. If any user is locked out, the account administrator can submit a support request in HPE GreenLake.

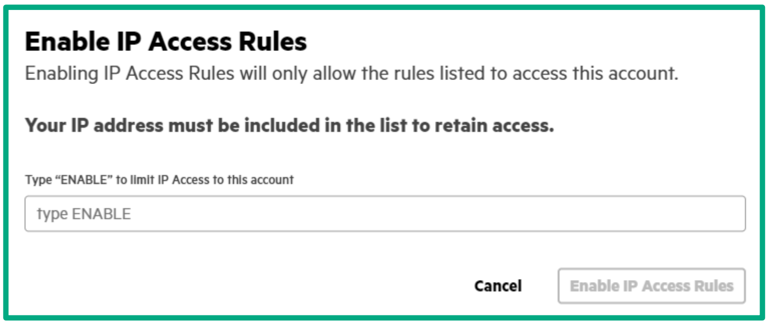

Enable IP Access Rules

Enabling IP Access Rules will only allow the rules listed to access this account.

Your IP address must be included in the list to retain access.

Type "ENABLE" to limit IP Access to this account

type ENABLE

Figure 12. HPE GreenLake Enabling IP Access Rules screen

Dual authorization

Dual authorization requires a secondary user to approve specific requests, such as delete resources or disable features. When enabled in the HPE Backup and Recovery service, any request to delete snapshots and backups must be approved by a secondary user with an administrator or equivalent role before it is executed.

Any request requiring dual authorization will be placed in a waiting state, until a second user can either approve or deny it. If a pending request is not explicitly approved or denied in seven days, it will be automatically denied.

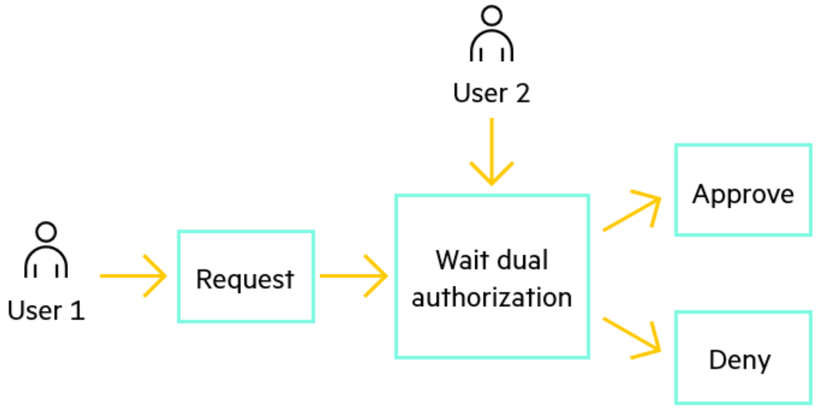

Figure 13. Dual authorization process

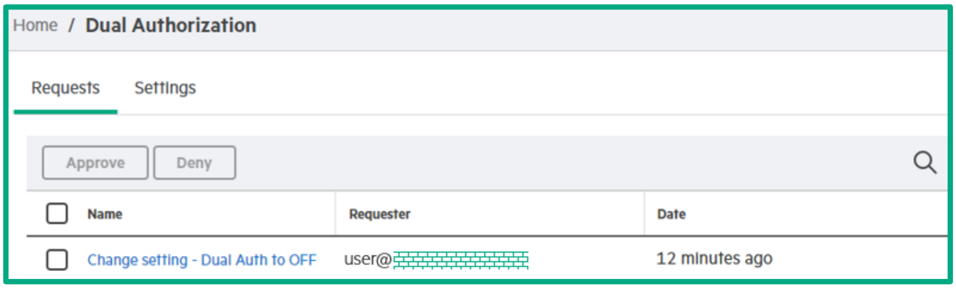

Home / **Dual Authorization**

☐	Name	Requester	Date
☐	Change setting - Dual Auth to OFF	user@[redacted]	12 minutes ago

Figure 14. Pending dual authorization requests in DSCC

Note
When dual authorization is enabled, users cannot approve their own requests. A secondary authorization is required to disable dual authorization in DSCC.

To approve or deny a request, the secondary administrator must explicitly confirm the action by entering `Yes`.

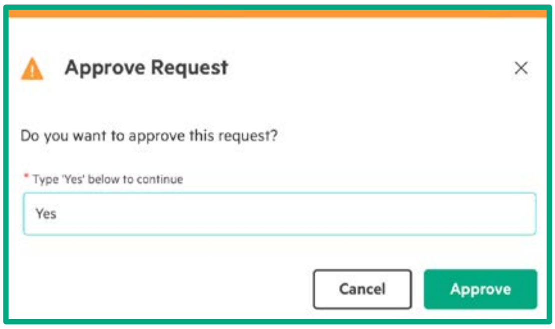

Approve Request

Do you want to approve this request?

* Type 'Yes' below to continue

Yes

Figure 15. Dual Authorization request approval confirmation

Audit logging

Audit logs are an essential security tool for providing records of all events and changes that occur within a system or environment.

The DSCC audit log service provides a comprehensive audit trail to assist in monitoring potentially sensitive data or systems for possible security breaches, vulnerabilities, or misuse of data. It also provides records that serve as evidence in cybersecurity attacks. In addition, audit logs may be considered business records that prove compliance with regulations and the law.

DSCC audit log access is customer-configurable to enable regular auditor checks for compliance purposes.

Table 9 provides a sample of audit log details.

Table 9. DSCC audit log example

Content	Description	Example
Who	The user who initiated the action	`John Doe`
Where	The source of the request	`10.10.10.10`
When	The timestamp when the request was made	`2022-09-23 12:48pm`
What	The action requested	`Login attempt`
Outcome	The result of the request	`Failed`

Note
Each DSCC application instance deploys a separate audit log.

API support

HPE GreenLake enables developers to make API calls on all array resources and capabilities of DSCC. The DSCC public API is specified based on the OpenAPI 3.0 format. It is an HTTPS end point, and it provides an X.509 server certificate for server authentication. Because the DSCC public API is presented as a single API, it has a single API root with a single version number covering all API resource groups. The major version number is included in the resource path (in the format `/api/v1/<resource group>`).

When newer versions of the API are released, both the new and old versions of API will be supported until the announcement of the deprecation for the old version of API. After it is deprecated, the older major version is frozen except for bug fixes. A deprecation header and sunset header is included in responses from the old API version to signal the change and the future withdrawal of the old API.

For more information about the FQDNs required for using the API, see the DSCC user and site public interfaces section of this paper.

Authentication

The DSCC public API uses the bearer token scheme in the HTTPS authorization header. A client application can access the API by using the access token, which is obtained from HPE GreenLake through the following process:

1. The user logs in to HPE GreenLake and navigates to the API section under **Manage Account**.
2. The user creates the credentials that will be associated with a specific instance of DSCC and will have an automatically generated client ID and client secret. The user must safely store this information.

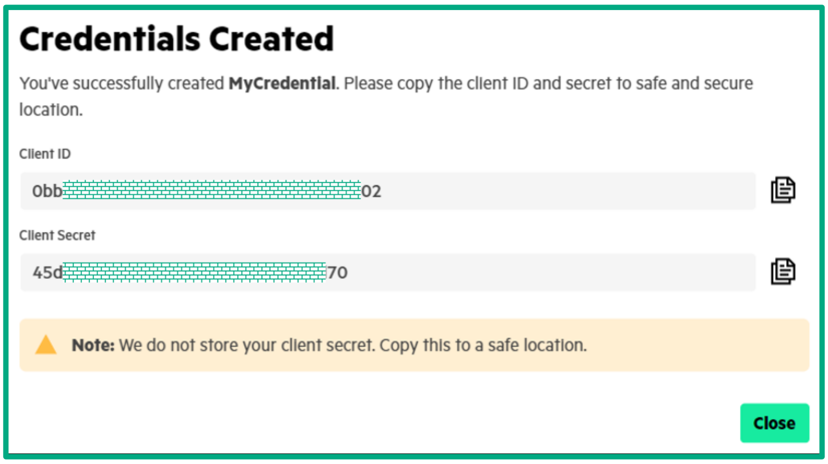

Credentials Created

You've successfully created **MyCredential**. Please copy the client ID and secret to safe and secure location.

Client ID

0bb[redacted]02

Client Secret

45d[redacted]70

Note: We do not store your client secret. Copy this to a safe location.

Figure 16. HPE GreenLake API Credentials Created screen

3. After the credential is created, an access token is generated.
4. The client application can then use the token by embedding it in the API request.

Authorization

The access token offers the client application the same level of authorization on resources as that of the user who created the access token. That is, through the access token, the client application inherits the same identity and authorization for the organization, role, and scope that are associated with the user who generated the access token.

Actions performed through the API are audited as if the user who owns the client ID and the secret was performing the action.

Access token security

The access token is valid for only 120 minutes, after which the customer must obtain a new access token.

In addition to the time limit of 120 minutes, the access token is invalidated under certain conditions:

- If the client ID or client secret is deleted.
- If the user account that was used to create the client ID and client secret is deleted.
- If the user revokes the access token.

A user can create a maximum of five API client credentials from a single HPE GreenLake user account. This rule helps to support multiple applications, and when one of the applications no longer requires access to the API, the user can revoke the associated access tokens, delete the client ID, or delete the client secret.

If the associated user's role changes so that the user loses access to the array, the API call to the array returns an unauthorized response.

Important

Resetting the HPE GreenLake user account password does not invalidate the access token.

To learn more about DSCC public REST API, refer to the HPE Developer Community.

Application credentials

HPE arrays and DSCC services require credentials to access and interact with other applications (for example, hypervisors) as part of data management workflows, such as backing up a virtual machine. Customers can provide application credentials of this nature to DSCC, where they are stored and protected in a secure credential vault. The secrets (keys) are encrypted for storage and transmission and are accessible only by the DSCC service that generates them. Each service has its own role and policies associated with the vault that allow it to get, list, create, and delete secrets. In addition, the only way the encrypted secrets are transmitted is over an HTTPS connection.

Compliance and certification standards

HPE ensures that HPE systems are secure, compliant, and certified. HPE GreenLake is focused on obtaining country-specific and industry-specific certifications and compliance:

- Processing of personal data by Hewlett Packard Enterprise in the context of the DSCC and HPE GreenLake complies with applicable privacy and data protection laws, including the General Data Protection Regulation (GDPR), and it is periodically reviewed to ensure continuous compliance.
- HPE GreenLake and DSCC, including the web portal and the connection between the on-premises device and the cloud, undergo continuous security penetration testing by third-party providers.

Fault tolerance of HPE GreenLake and DSCC

Because customers use DSCC to manage their storage fleet, they rightly expect it to be highly available to enable them to perform management tasks as needed. To protect this availability, DSCC instances are distributed in multiple availability zones in each region. The DSCC instances are configured so that if a component in an availability zone fails, it is brought up seamlessly in the same or another availability zone within the same region, and the on-premises devices can automatically reconnect. These availability zones are interconnected with high-bandwidth, low-latency networking that has full redundancy between availability zones.

In addition, the infrastructure is regularly backed up through automated processes and can be redeployed quickly in case of a major failure.

Summary

DSCC provides a secure, common, and consistent customer experience for all services related to storage, system management, data management and data protection engineered by HPE Storage, fully integrated with HPE GreenLake features and services. DSCC offers several features the customer can use to improve the security of their environment, while enabling a cloud-native unified management experience.

The customer's organization must have confidence that their infrastructure components, such as their data storage fleet, can meet and maintain their organization's security policies. Security involves all aspects of the IT infrastructure—from the smallest component up through the collection of computing devices in the data center, their connectivity, and the ways all these components are managed.

Hewlett Packard Enterprise understands that organizations do not want to allow unbridled access to their environment. All devices are configured by using a least-privilege level of access, and all access is controlled by the customer. Therefore, customers determine which service personnel can have access to their devices and when that access can occur. With hardware and software based on these principles, the security of their data is always under their control.

Additional links

DSCC technical white paper
https://www.hpe.com/psnow/doc/a00124553enw

HPE GreenLake edge-to-cloud platform user guide
https://support.hpe.com/hpesc/public/docDisplay?docId=a00120892en_us

Getting started with DSCC
https://infosight.hpe.com/InfoSight/media/cms/active/public/pubs_Data_Services_Cloud_Console_Getting_started_Guide_60x.whz/index.html#ncl1622811349917.html

HPE Storage welcome center
https://infosight.hpe.com/welcomecenter

HPE Alletra 9000 storage simplified setup with DSCC
https://www.hpe.com/psnow/doc/a00121299enw

DSCC REST API blogs in HPE developer community
https://developer.hpe.com/platform/data-services-cloud-console/home/

HPE Remote Device Access (RDA) security white paper
https://support.hpe.com/hpesc/public/docDisplay?docId=a00006791en_us

HPE global privacy policy
https://www.hpe.com/us/en/privacy/master-policy.html

Learn more at

hpe.com/storage